

Remarks

Support for the amended claims

Further experience with the redirection mechanisms disclosed in the present patent application
5 has resulted in the realization that using the redirection mechanisms as disclosed in the present
patent application to deal with the lack of an object specified in a request received in a database
system belonging to a distributed database system is only one species of a generic invention,
namely using the redirection mechanisms in any circumstance where it becomes apparent during
execution of the request that "the request is preferably executed at least in part in another
10 database system" (claim 5, line 7) of the distributed database system. The disclosure in the
application as filed of the species that deals with lacking objects of course supports the generic
claim. Support is also provided by the description of what may constitute a miss at page 13,
lines 8-14 of the application as filed:

15 There are many reasons why a miss may occur. Perhaps the most common is that
an object referred to in the request is not present in database system 203; another
is that the request updates an object and the object is read only in database system
203 but updateable in database system 229; a third is that the request uses
20 information about an object that may be different in database system 203 and in
database system 229; a fourth is that the user making the request has privileges in
database system 203 that prohibit the request and privileges in database system
229 that allow it.

Examiner will immediately see that the application as filed supports dependent method claims 6-
25 9 and dependent Beauregard method claims 10-14. Apparatus claims 15-19 and dependent
Beauregard apparatus claims 20-24 are supported in the application as filed at least by FIG. 3
and the discussion of FIG. 3, which explains at page 17, line 14 through page 18, line 14 how
semantic analyzer 305 detects a reference to an object that is not in database 203 and how
redirectors 307 and 317 redirect the request.

Patentability of the claims over the references

In the following, Applicants will first describe what Applicants are claiming and will then
demonstrate that the references fail to disclose the limitations of Applicants' claims either
individually or in combination.

What Applicants are claiming

Applicants' claim 5 broadly sets forth Applicants' invention as a "method employed in a distributed database system . . . for responding to a request":

1 **5. (new)** A method employed in a distributed database system that includes a
2 plurality of database systems for responding to a request received in a database
3 system of the plurality,
4 the method comprising the steps performed during execution of the request in the
5 database system of:
6 determining that the request is preferably executed at least in part in
7 another database system of the plurality; and
8 redirecting the execution of at least the part of the request to the other
9 database system.

10 Beginning with definitions of important terms, as may be seen from FIG. 1, the "distributed
database system" includes "a plurality of database systems". It is clear from the discussion
beginning at page 9, col. 8 that a "database system" is a system that can respond to requests that
include queries written in a query language such as SQL. It is further important to claim 1 that
15 the steps of the method are performed "during execution of the request in a database system of
the plurality". Finally, the method steps involve "another database system of the plurality", that
is, another database system belonging to the distributed database system.

*What the references disclose*Bodnar and Klots

20 Bodnar and Klots contain no disclosure concerning distributed database systems as that term is
used in Applicants' claims. Bodnar and Klots disclose caches in local nodes of a system. The
caches contain information that is obtained from a database system accessible to the nodes, but
25 the caches are not themselves database systems. In the case of Bodnar, what is disclosed is well
set forth in the Abstract:

30 A system and method for caching frequently called numbers so that call delivery
time is shortened and the signaling network and service nodes are not overloaded
by telephone number queries. In a distributed control switching system, each
processor that performs digit analysis and has an area of memory. This area of
memory is used as a cache to store directory number translation information.
When a call arrives for a specified directory number (DN), if, after standard
translation, the call cannot be normally routed, then the cache is consulted to
determine whether the destination for that DN is in the cache. If it is, then the call
35 is routed to the destination associated with the DN. If, however, the DN is not in

the local cache, then a query is sent to one or more other units in the switch which may contain such caches in order to determine if these caches contain the destination. This "distributed" query is still faster than querying a centralized database because all data queries throughout the switch are carried out in parallel.

5

Bodnar thus discloses a system in which a cache in a module of a switch is used to store information needed to translate telephone numbers. The source of the information is a database that is accessible via the telephone network. If the needed information is not in a module's cache, the module queries other modules to determine whether their caches have the information.

10

If none does, the source database is queried. There is no indication anywhere in Bodnar that the caches are "database systems" as that term is used in Applicants' claims, and there is no technical reason why they would be. Bodnar thus does not disclose the "distributed database system" of Applicants' claim and consequently cannot, as a matter of simple logic, disclose redirection from one database system in the distributed database system to another database system in the distributed database system in the manner set forth in the claim.

15

The same is true with Klots. As shown in FIG. 1b and set forth at col. 2, line 6-col. 3, line 2, the system disclosed in Klots has a single database which is accessed by a plurality of nodes. Each node includes a cache, and the problem addressed in Klots is dealing with the fact that the current version of a data item required by one of the nodes may be in a cache in another of the nodes. This is clearly set forth in Klots' *Abstract*:

20

25

30

35

The method involves determining whether a current version of the data item resides in a cache that is remote relative to the first database server. If the current version of the data item does not reside in a cache that is remote relative to the first database server, then a first derivation mechanism that is local to the first database server derives the particular version of the data item from the current version of the data item. If the current version of the data item resides in a cache that is remote relative to the first database server, then a second derivation mechanism that is local to the cache derives the particular version of the data item from the current version of the data item. Once built, the particular version of the data item sent to the first database server. Determining whether a current version of the data item resides in a cache that is remote relative to the first database server may be performed by determining whether a process currently holds a write lock on the data item, and if a process currently holds a write lock on the data item, then determining whether the process is remote relative to the first database server.

As was the case with Bodnar, there is no indication anywhere in Klots that the caches are "database systems" as that term is used in Applicants' claims, and there is no technical reason

why they would be. Klots thus does not disclose the "distributed database system" of Applicants' claim and also cannot, as a matter of simple logic, disclose redirection from one database system in the distributed database system to another database system in the distributed database system in the manner set forth in the claim.

5

Cusson

Cusson is a CIP of USSN 09/294,656 and contains the entire *Detailed Description* from that application. USSN 09/294,656 is indicated as a "related application" in the present application's *Cross references to related applications*. A PCT application corresponding to USSN 09/294,656 was published as WO 00/63800. USSN 09/294,656 and Cusson describe a distributed database system in which a first database system serves as a cache for a second database system and queries which cannot be resolved in the first database system are redirected to the second database system. The present patent application sets forth the differences between its distributed database system and that of Cusson as follows at page 13, lines 16-24:

15 Database system 201 differs from the system disclosed in WO 00/63800, Cusson, et al., *Web servers with queryable dynamic caches*, in that the redirection is done *within* database system 203, not in the data access layer. It should further be pointed out here that the techniques of database system 201 only require that all of the references in the request received in database system 203 be resolvable in one or another of a set of database systems that either themselves redirect SQL requests or are reachable by redirection from database systems that do, with all of the databases in the set being reachable either directly or indirectly by a redirection that begins from database system 203.

25 The foregoing difference is clearly set forth in claim 5, which requires that the steps of the method be "performed during execution of the request in a database system of the plurality".

That redirection is not done *within* Cusson's database system is apparent from FIG. 2 of Cusson and the discussion of that figure which begins at col. 5, line 26 of the reference. FIG. 2 shows a network server 203(i) which includes a cache database 236. As shown by the description of the operation of cache database 236 at col.6, lines 30-40, the presence of query engine 221 in cache database 236, and the fact that cache database 236 responds to cache queries 245 and provides cache results 247, cache database 236 is a database system. Redirection of requests from cache database 236 to source database 241 is, however, not done "during execution of the request" in cache database 236, but rather by data access layer 253, which is completely separate from cache database 236. Moreover, the determination whether to redirect is not made in cache

database 236, but rather in data access interface 219, as is explained at col. 6, lines 1-29. Examiner bases his rejection of original claim 1 on col. 4, lines 31-35 of Cusson. The "miss table" mentioned there is part of data access interface 304, an embodiment of data access interface 219. See in this regard FIG. 3 of Cusson and col. 12, lines 51-52. Thus, Cusson does not disclose the limitation that the steps of the method of claim 1 be "performed during execution of the request in a database system".

Detailed discussion of the rejections

It will be immediately apparent that because Klots and Bodnar disclose nothing whatever about distributed database systems, and because redirection in those references takes place not between database systems in a distributed database system, but rather between nodes that have caches that are not database systems, these references cannot individually anticipate claim 1 or be combined with any of the other cited references to render claim 1 obvious.

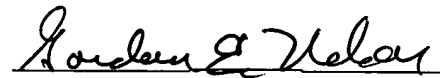
It will also be immediately apparent that because Cusson does not perform the steps of "determining that the request is preferably executed at least in part in another database system of the plurality" and "redirecting the execution of at least the part of the request to the other database system" "during execution of the request in a database system of the plurality", it cannot anticipate claim 1 or be combined with any of the other cited references to render claim 1 obvious.

It will further be apparent to Examiner that apparatus claim 15 is neither anticipated nor rendered obvious by the references for the same reasons that claim 1 is neither anticipated nor rendered obvious by them. Since the two independent claims in the application are patentable over the references, so are all of the dependent claims. It should however be further pointed out claims 7-9, 12-14, 17-19, and 22-24 are also patentable in their own rights over the references. These claims address the techniques used to do the redirection within the database system, and none of the references provide any disclosure whatever of what happens within any database systems disclosed in the references. They therefore do not disclose any of the added limitations of claims 7-9, 12-14, 16-18, and 21-23 and consequently cannot be combined with other references to render these claims obvious.

Conclusion

In the foregoing, Applicants have shown that their new generic claims are fully supported by the application as filed and have shown that the generic claims are patentable over the references cited by Examiner. That being the case, Applicants have completely satisfied the requirements of 37 C.F.R. 1.111(b) and respectfully request that Examiner continue with the examination as provided by 37 C.F.R. 1.111(a) and allow the claims as amended. No additional fees are believed to be required for this amendment. Should any be, please charge them to deposit account number 501315.

Respectfully submitted,



Attorney of record,

Gordon E. Nelson

57 Central St., P.O. Box 782

Rowley, MA, 01969,

Registration number 30,093

Voice: (978) 948-7632

Fax: (866) 723-0359

12/10/2003

Date

Certificate of Mailing

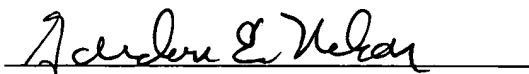
I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to:

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

on 12/10/2003

(Date)

Gordon E. Nelson, #30,093



(Signature)